Marginal Propensity to Consume of Consumption and Savings Targeters

Pavel Koval *
University of Minnesota

February 2025

Abstract

This paper provides an explanation for the conflicting evidence on the relation between marginal propensity to consume and liquid assets. I study marginal propensity to consume (MPC) out of liquid assets in an environment where agents infrequently update their financial plans and choose between consumption and savings as a decision variable in the utility maximization problem. I show that the average MPC consists of the average slope of policy functions in the economy and additional component resulting from agents switching between the decision variables. I assess the component in a dynamical consumption-savings model with income fluctuations, and demonstrate that the presence of this component can explain conflicting evidence on the negative relation between MPC and liquid assets.

Keywords: marginal propensity to consume, consumption-savings model, inattention, liquid assets

*Email: koval084@umn.edu. I thank Filip Matejka, Ricardo Reis, V.V. Chari, Wouter DenHaan, Ben Moll, Ethan Ilzetzki, Marek Kapicka, Ctirad Slavik and participants of several seminars at CERGE-EI and LSE for useful comments and suggestions. This project has received funding from the European Research Council under the European Union's Horizon 2020 research and innovation programme (grant agreements No. 101002898).

Introduction

Standard consumption-savings models with income fluctuations predict that consumption policy functions are concave in liquid assets (Zeldes, 1989; Carroll and Kimball, 1996), and thus, marginal propensity to consume (MPC) is decreasing in liquid assets. Recent empirical evidence (Lewis et al., 2019) find that the relation between MPC and liquid assets is, on average, positive. This paper shows how standard theory can be reconciled with this empirical evidence without sacrificing on the concave relation between consumption and assets. I develop an income fluctuation model, in which agents choose between consumption and savings as a decision variable in their utility maximization problem and update their plans infrequently. A small change in assets induces a fraction of population to switch from savings to consumption as their decision variable, which results in an additional component in the aggregate MPC. This component is negative and concentrated in the low levels of assets, driving average MPC in the low levels of assets down. As a result, the relation between aggregate MPC and liquid assets becomes negative when the component is sufficiently large. I assess the size of the component in a dynamical consumption-savings model with income fluctuations and show that my explanation is plausible.

As Kaplan and Violante (2022) laid out, standard macroeconomic models predict a negative relation between MPC and the level of liquid assets for several reasons. First, some poor agents may face binding credit constraints and thus consume the majority of any extra liquidity they receive. Second, uninsurable risk, combined with a consumption smoothing motive from concave utility, creates two forces that increase the MPC due to prudence (Carroll and Kimball, 1996) or an occasionally binding liquidity constraint (Zeldes, 1989; Deaton, 1991; Carroll and Kimball, 2001). These precautionary saving motives are the most pronounced at the low levels of liquid assets but disappear as liquid assets grow sufficiently large. As a result, precautionary motives lower the level of consumption relative to the certainty equivalent case, which is linear in liquid assets but does so more at low liquid assets levels. The consumption function is therefore concave in liquid assets.

However, empirical evidence doesn't unequivocally support this theoretical prediction. Recent paper by (Lewis et al., 2019) find a positive significant relation between MPC and the level of income/liquid assets. This finding does not stand alone (Shapiro and Slemrod, 2009; Broda and Parker, 2014; Kueng, 2018), suggesting the need for theoretical explanation [1]. Lewis et al. (2019) conjecture that their results could be explained by the dependence of MPC on the relative size of the payment, theories of near-rationality or mental accounting (Boutros, 2022), and time-varying consumption thresholds (Miranda-Pinto et al., 2020).

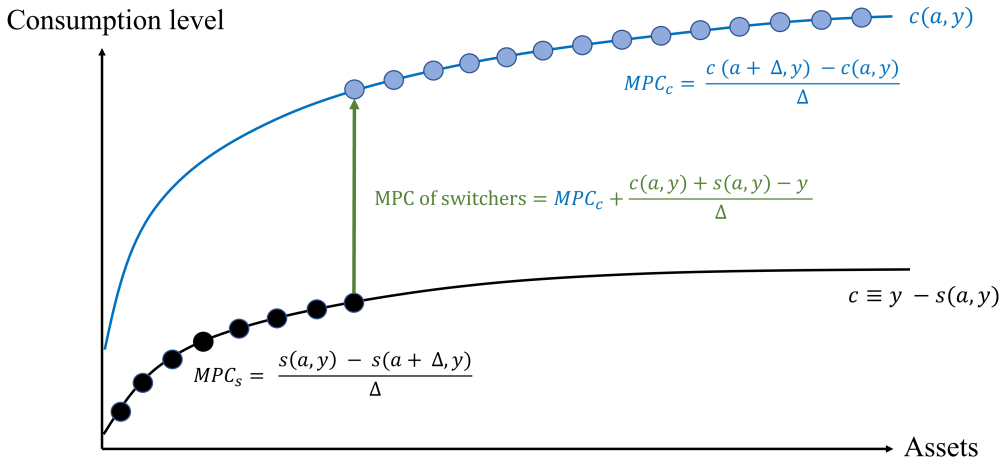
[1] However, some papers (Jappelli and Pistaferri, 2014; Fagereng et al., 2021) find the relation to be negative, in line with the standard theoretical view. I explain how someone could find negative relation further in the text

This paper shows that the average MPC could positively depend on liquid assets even if each agent in the economy has a concave consumption policy function, in line with the standard theoretical view. As I demonstrate below, in my model, average MPC consists of the average slope of policy functions *and* additional component, which is *negative* and *more pronounced in the low levels of liquid assets*. As a result, the overall relation can be zero or even reverse sign if the negative component in the low levels of assets is sufficiently large.

This negative component arises due to the difference in economic outcomes from using consumption or savings as a decision variable in the utility maximization problem. Agents infrequently update their plans in the spirit of Reis (2006) and can choose between using consumption or savings as a decision variable. I show that the choice is determined by agent's income state and liquid assets. Consequently, when the agent experiences an increase in her liquid assets, she might want to switch between decision variables. As I show further, there are more switchers in the low levels of liquid assets and switching behavior induces a negative component of average MPC. Therefore, one could observe a positive (or zero) correlation between MPC and liquid assets level in the data even though agents have concave consumption policy functions, in line with the standard income fluctuations model.

To illustrate the argument, in Figure 1 I depict two policy functions that determine consumption levels for any asset level a, given income y. The blue curve depicts the consumption policy function $c(a, y)$, which determines the optimal amount of consumption given assets a and income y. The black curve depicts the consumption level determined by the savings policy function $s(a, y)$, which prescribes the optimal amount of savings given resources. The consumption level is defined as $c \equiv y - s(a, y)$. In the absence of any frictions in consumption behavior, there is no discrepancy between $c(a, y)$ and $c \equiv y - s(a, y)$. Let us assume there is a friction, which induces the discrepancy, and an agent is allowed to choose between using either of the policy functions. Let us also assume that there is a threshold $\hat{a}$ *below* which it is optimal to use the *savings* policy function and *above* which it is optimal to use the *consumption* policy function. Then, MPC below the threshold $\hat{a}$ is defined by the slope of the consumption policy function and MPC above the threshold $\hat{a}$ is defined by the slope of the savings policy function. Crucially, if the agent's asset level is slightly below the threshold $\hat{a}$, it might be optimal for her to *switch* from the consumption policy function to the savings policy function after receiving additional assets Δ. Her MPC is then defined by the difference between the two policy functions and does not coincide with the slope of either of the functions. The expression $c(a, y) + s(a, y) - y$ is determined by the severity of the behavioral friction. Once one aggregates over different MPCs in the population, the wedge between the average MPC and the average slope of policy functions occurs, and the component $c(a, y) + s(a, y) - y$ is exactly the source of the wedge. Even though the wedge arises because of switchers only, it is non-negligible as I show further.

Figure 1: Optimal Plans



To develop a model that allows for the discrepancy between $c(a, y)$ and $c \equiv y - s(a, y)$, I use several empirical facts on people's planning behavior. According to the Survey of Consumer Expectations (SCE), some people plan their spending by trying to "keep spending within a certain limit" while others "focus more on trying to save regular amounts of money". The former approach suggests a focus on the consumption level and the latter approach clearly indicates a focus on the savings level while making economic decisions. Throughout this paper, I call the agents who follow the former approach "consumption targeters" and the latter approach "savings targeters". The data shows that consumption targeters and savings targeters are different in multiple dimensions; in particular, savings targeters accumulate fewer assets but have a higher annual income than consumption targeters. In the model, I define consumption targeters as those who use consumption as a decision variable (blue curve $c(a, y)$ on Figure 1) and savings targeters as those who use savings as a decision variable (black curve $y - s(a, y)$ on Figure 1). Also, Andre et al. (2024) conducted a large-scale survey on consumption and savings behavior and show that 68% of the households follow simple quick-fix consumption rules that either fully consume or fully save, in line with the definition of consumption targeters and savings targeters in this paper.

Furthermore, empirical evidence suggests that people tend to plan their spending infrequently. This could originate from the limited mental capacity of agents resulting and their inattentiveness to relevant economic outcomes (Pope, 2009; Chetty et al., 2009; Mondria and Quintana-Domeque, 2013; Bartoš et al., 2016). Further, the recent evidence on spending budgeting (Zhang et al., 2022) documents that people who create spending plans update them infrequently. Most of the respondents in a survey conducted by Zhang et al. (2022) assess their spending budgets only once or twice a month, while their consumption decisions are much more frequent. The assumption that people infrequently update their plans drives the difference in allocations between consumption and savings targeters in my model.

The model predicts that the choice of becoming either a consumption or savings targeter is defined by the amount of assets and income and is influenced by the trade-off between certainty in assets or consumption. By explicitly choosing a savings level, a savings targeter ensures certainty in assets but experiences uncertainty in her consumption level. She commits to saving a certain amount today and tomorrow for any income realization and thus consumes tomorrow more or less (than expected) depending on income realization. Conversely, a consumption targeter guarantees a certain consumption level by absorbing income uncertainty with her assets. She commits to consuming a certain amount and saves more or less depending on income realization. The model predicts that an agent decides to become a savings targeter if she has relatively low accumulated assets and a high income. Intuitively, this happens because, by choosing the amount of savings, she guarantees more certainty in her future assets, which is valuable when her assets are low. However, if the income level today is low, it will, most probably, be low tomorrow due to income persistence, and thus the consumption uncertainty becomes sensitive. Therefore, the agent prefers to become a consumption targeter and to guarantee a certain amount of consumption if her income is relatively low. The data confirm the model predictions.

The intuition for the main result of the paper is the following. The difference between two types of behavior is more pronounces at the low level of assets (for a given income) because uncertainty in assets induced by being a consumption targeter is more hurtful when the utility is lower. As a result, switching between types resulting from the change in assets is predominantly concentrated in the low levels of assets and so is the component $c(a, y) + s(a, y) - y$. This component is negative because the consumption level of a consumption targeter $c(a, y)$ is lower then of a savings targeter $c \equiv y - s(a, y)$, as the former effectively has higher uncertainty in her future assets, and thus her precautionary motive is stronger. Therefore, the aggregate MPC can have zero or negative correlation with liquid assets if $c(a, y) + s(a, y) - y$ is sufficiently large.

This paper aims to complement the rich body of literature that models inattentive consumption behavior using multiple different approaches, e.g. rational inattention (Sims, 2003; Tutino, 2013; Kőszegi and Matějka, 2020), salience (Bordalo et al., 2013), focusing (Kőszegi and Szeidl, 2013), and sparsity (Gabaix, 2014). This paper follows the approach of modeling inattentiveness as infrequent updating of plans (Gabaix and Laibson, 2002; Reis, 2006). Almost all the papers that model consumption behavior implicitly assume that consumption and savings decisions are interchangeable; the few exceptions are Reis (2006) and Gabaix (2016). The closest related paper in this direction by Reis (2006), who also modeled the decision of becoming a consumption or savings targeter, has a crucial distinct feature. In my model, the decision is state-dependent (depends on the asset and income levels) while Reis's (2006) is not, so his model could not account for those who change their type because parameters driving the choice were assumed to be unchangeable. As I show, the existence of

switchers is supported by data and is crucial for the aggregate response of consumption to fiscal stimuli. I develop a novel model of consumption behavior that explains the difference in outcomes from the usage of consumption or savings as a decision variable and that matches micro-level data.

The remainder of this paper is organized as follows. Firstly, I present the toy model, which explains the driving mechanism of the type choice, and then support the main prediction of the model with empirical evidence. I then explore the average MPC in the economy of consumption targeters and savings targeters and show that there is a wedge between the average MPC and the slope of policy functions. Further, I develop and solve an infinite-horizon model to calibrate the wedge induced by the switchers.

1 Toy model and empirical evidence

Toy model of consumption and savings targeters

This toy model serves to explain the trade-off between becoming a consumption targeter [CT] or a savings targeter [ST]. In the model, an agent lives for three periods and observes her income in the first period y_1 but does not know her income in the second y_2 and third y_3 periods. Income realizations are drawn from the same distribution with mean μ and variance σ^2. I assume that income realization in the first period is the mean value $y_1 = \mu$. I define savings as $s \equiv y - c$. The agent creates plans for all three periods and does not update them as time goes by. After observing income in the first period μ, the agent chooses between consumption and savings plans by comparing expected utilities from optimally chosen values.

The agent might create a consumption plan (use consumption as a decision variable):

$$W^c(a, y_1) = \max_{c=(c_1,c_2)} u(c_1) + \beta u(c_2) + \beta^2 E[u(a + y_1 + y_2 + y_3 - c_1 - c_2)] \quad s.t. \tag{1}$$

$$c_1 > 0,\ c_2 > 0,\ c_1 + c_2 \leq a + 2\underline{y} + y_1$$

Following this plan, for any income realizations, the agent consumes a *certain* amount in the *first two periods*. In the third period, she faces uncertainty in her accumulated assets $a' = a + \mu - c_1 + y_2 - c_2$ and income realization y_3. I call the agent who follows this type of behavior a *consumption targeter [CT]*. The amount $\underline{y}$ in the constraints is the least possible realization of income. The agent cannot plan to consume more than her total wealth once the least possible income realizes in the first two periods.

The agent might create a savings plan (use savings as a decision variable):

$$W^s(a, y_1) = \max_{s=(s_1,s_2)} u(y_1 - s_1) + E[\beta u(y_2 - s_2) + \beta^2 u(a + y_3 + s_1 + s_2)] \quad s.t. \tag{2}$$

$$s_1 \leq y_1,\ s_2 \leq \underline{y},\ a + \underline{y} + s_1 + s_2 > 0$$

Following this plan, the agent *saves* a certain amount in the first two periods, keeping her consumption in the second and third periods uncertain. The agent who follows this type of behavior I call a *savings targeter [ST]*. In the third period, she knows her accumulated assets $a' = a + s_1 + s_2$ with certainty and faces uncertainty only in income realization y_3.

Ultimately, the agent chooses between types by comparing expected utilities:

$$W = max\{W^c, W^s\}. \tag{3}$$

The crucial difference between the types is in their attitudes toward uncertainty. Firstly, a consumption targeter postpones uncertainty from the second period to the third period while a savings targeter faces uncertainty in both the second and third periods. Secondly, a savings targeter "splits" the uncertainty between the second and third periods, while a consumption targeter "bundles" the uncertainty from the second and third periods together. Thus, a consumption targeter faces a tighter constraint as she accounts for the fact that income could take the least possible amount *twice*.

For the quadratic utility, the model can be solved in the closed form. Even though the quadratic utility is known to be unrealistic, it perfectly depicts the difference between the CT's and the ST's attitudes towards uncertainty.

If none of the constraints binds, it is always better to be CT than ST and the difference between the two value functions is $W^c - W^s = (\beta - \beta^2)\sigma^2$. This difference arises precisely because CT postpones uncertainty to the third period. As the agent is impatient $\beta < 1$ and risk-averse, the difference is positive. This result is expressed in Lemma 1 (for the proof see Appendix A).

Lemma 1: *Under quadratic utility, unconstrained problems (1) [CT] and (2) [ST] are equivalent in terms of providing the same level of maximized utility and expected allocations if $\beta = 1$; unconstrained problem (1) [CT] is strictly preferable to unconstrained problem (2) [ST] if $\beta < 1$.*

When decreasing a or increasing $\mu = y_1$ (namely, I vary both the mean of the income distribution and first-period consumption), the non-negativity constraint of a consumption targeter in the third period ($a + 2\underline{y} + y_1 - c_1^* - c_2^* \geq 0$) always binds before any non-negativity constraint of a savings targeter. This happens precisely because CT bundles the uncertainty, so her constraint is tighter than that of ST who splits the uncertainty between periods. Plugging in the optimal consumption levels, one could obtain the following expression for the constraint $a \geq (2\beta^2 + 2\beta - 1)\mu - 2\underline{y}(1 + \beta + \beta^2)$, which is more slack for *higher* assets

a and *lower* income mean μ. Under quadratic utility, the agent optimally consumes based on average income realizations μ in each period, while the constraint involves wealth for the two worst realizations of income $2\underline{y}$. As a result, for high enough μ it will be optimal to consume more than the worst realization of the third-period wealth and the constraint will bind.

Consequently, for lower enough assets a or high enough income state μ, the difference between constrained and unconstrained value functions of CT will become lower than the constant $(\beta - \beta^2)\sigma^2$, and the agent will optimally prefer to become ST. This result is expressed in Proposition 1 (for the proof see Appendix B).

Proposition 1: *Under quadratic utility, there is a line in a state space of assets and income with a positive slope above which it is optimal to be ST and below which it is optimal to be a CT.*
For $\beta = 1$, the expression $\mu - 2\underline{y} - \frac{a}{3}$ serves as sufficient statistics for the choice of a type.

To summarize, without the non-negativity constraints, it is always better to be CT than ST because both types dislike uncertainty and CT experiences it further in the future. However, CT bundles uncertainty from the second and third periods together while ST splits it between periods, so the non-negativity constraint of CT always binds first. The constraint is tighter for lower assets a and higher income state μ, and thus for low enough a or high enough μ it is preferable to be ST. This result is presented in Figure 2.

Figure 2: Optimal Plans

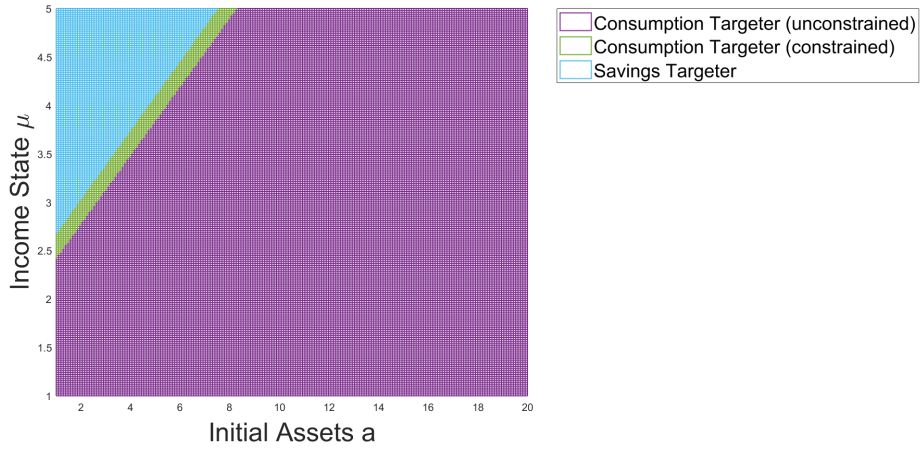


Figure 2 depicts the decision of an agent for different values of income state μ and asset level a, which is determined by equation (3). In the blue area, the savings plan provides higher expected utility than the consumption plan, and thus the agent becomes a savings targeter.

In the purple area, a consumption plan provides higher expected utility than a savings plan, and thus the agent becomes a consumption targeter. In the green area, a constrained consumption plan is preferable to a savings plan. One could see that savings targeters are those with higher first-period income for a given level of assets, while consumption targeters are those with a higher level pf assets for a given first-period income. The line between blue and green areas is exactly the one described in Proposition 1.

Under quadratic utility, an agent does not exhibit prudence and thus does not have a precautionary motive (Kimball, 1990). This makes ST a passive benchmark for the comparison with CT, because, in this framework, uncertainty matters only when a constraint binds. For more general utility functions discussed in Sections 2 and 3, CT and ST also become different in their response to the uncertainty they face.

The observational implication of Proposition 1 is that for any given asset level one will observe savings targeters to have, on average, higher income than consumption targeters. Furthermore, for any given income level one will observe that savings targeters have, on average, lower amounts of assets than consumption targeters. I test this implication using data on households' financial behavior.

Empirical evidence

The main data source I use to grasp information about consumption targeters and savings targeters is the SCE household finance survey. This survey includes a direct question about the type of planning behavior of a household. If a household member answers that she plans financial behavior by trying to "keep her spending within a certain limit", I classify her as a consumption targeter [CT]. If she does it by trying to "focus more on trying to save regular amounts of money" or "pay off regular amounts of debt", I classify her as a savings targeter [ST]. Among those who try to budget, 62% of respondents are consumption targeters and 38% are savings targeters. Importantly, agents state their type explicitly, and thus I could classify them without relying on any specific assumption on the income process and expectation formation.

Table 1 provides summary statistics of agents' characteristics by their type. On average, consumption targeters have lower income but higher assets than savings targeters, which is in line with the main prediction from the model. In addition, consumption targeters are, on average, older and work less in full-time jobs, suggesting the life-cycle dimension of the choice between types.

Table 1: Summary statistics

	Consumer type	Saver type	Δ
Number of units	763	469	
Annual Income (thousand dollars)	65.0	77.5	-12.5^{***}
Total Assets (thousand dollars)	184.8	141.5	43.3^{**}
Age (years)	52.6	47.5	5.1^{***}
Working full-time (share)	0.50	0.69	-0.19^{**}

Importantly, the table depicts only unconditional means for the types, but the model provides predictions about the conditional relationship between income and assets. Namely, consumption targeters should have higher assets for a given income level, while savings targeters should have higher income for a given level of assets. To take these facts into account, I run the following regression.

$$log(assets_i) = \alpha_1 + \alpha_2 log(income_i) + \alpha_3 consumer_i + \alpha_4 X_i + \xi_i, \tag{4}$$

$$log(income_i) = \gamma_1 + \gamma_2 log(assets_i) + \gamma_3 consumer_i + \gamma_4 X_i + \xi_i, \tag{5}$$

where $assets_i$ is a total assets level (sum of assets in defined contribution plans, self-owned retirement accounts, and savings-investments accounts), $income_i$ is an annual income level, $consumer_i$ is a dummy for an individual to be a consumption targeter, X_i is a set of controls that includes age, employment, and education. The main coefficients of interest are α_3 and γ_3 and the theoretical models predict them to be of different signs $\alpha_3 > 0$ and $\gamma_3 < 0$.

The result depicted in Table 2 is in line with the theoretical predictions. For a given income level, consumption targeters have on average 41.3 p.p. higher assets than savings targeters. Conversely, consumption targeters have on average 16.8 p.p. lower annual income than savings targeters with the same assets level.

Even though the data provides a useful identification of types, it has several crucial limitations. The main one is that income and assets are documented on annual frequency, while Zhang et al. (2022) argue that most people update their plans on a monthly frequency. Thus, the data lack variation in the type choice related to a variation in income and assets.

The data also provides only limited evidence on switchers. For example, I cannot assess whether switchers from a consumption targeter type to a savings targeter type experienced a positive income change as the theoretical model predicts. This is because the dataset on income and assets is of annual frequency and none of the respondents stays in a sample long enough to experience both changes in type and change in income (assets). Therefore, I can provide information only on the share of switchers, which is around 20% of those who stay

Table 2: Empirical evidence

	Dependent variable:	
	log(assets)	log(income)
	(1)	(2)
Consumer type	0.413^{***}	-0.168^{***}
	(0.141)	(0.056)
log(income)	1.016^{***}	
	(0.071)	
log(assets)		0.161^{***}
		(0.011)
Observations	1,050	1,050
Controls & Constant	Yes	Yes
Note:	*p<0.1; **p<0.05; ***p<0.01	

in a sample at least for two quarters.

2 General framework and Marginal Propensity to Consume

In this section, I describe the general framework that I use to analyze the decision of becoming a consumption or a savings targeter and the consumption behavior of these two types.

Assume there is a continuum of agents of measure one in the economy. The agents are ex-ante heterogenous in the initial assets $a \in A$ and income $y \in Y$. Assume also there are two value functions $V^c(a, y)$ and $V^s(a, y)$, which are maximized lifetime utility flows. They both are functions of assets a and income y, and differ by the decision variable that is used to maximize the lifetime utility flow. An agent is a *consumption targeter* if she uses consumption as a decision variable $c^*(a, y)$, and her savings are defined by the residual $s \equiv y - c^*$. An agent is a *savings targeter* if she uses savings as a decision variable $s^*(a, y)$, and her consumption is defined by the residual $c \equiv y - s^*$. The choice between decision variables $x^* = \{c^*, s^*\}$ is made by comparing two value functions:

$$V(a, y) = max\{V^c(a, y), V^s(a, y)\} \tag{6}$$

The choice could be one-time (as in the toy model) or repeated multiple times (as in the

infinite-horizon model presented in Section 3).

Assume also $V^c(a, y)$ and $V^s(a, y)$ are such that there is a unique continuous bijective mapping from asset space A into income space Y, $f(a)$, such that $V^c(a, f(a)) = V^s(a, f(a))$. By taking the full derivative and re-grouping, I obtain the following expression for the slope of the function $f(a)$:

$$\frac{\partial f(a)}{\partial a} = \frac{V_1^c(a, f(a)) - V_1^s(a, f(a))}{V_2^s(a, f(a)) - V_2^c(a, f(a))}, \tag{7}$$

where V_1 denotes the derivative of a value function with respect to the first argument (assets) and V_2 denotes the derivative of a value function with respect to the second argument (income).

Lemma 2: *For all pairs* $(a, f(a))$ *such that* $V^c(a, f(a)) = V^s(a, f(a))$*, the slope* $\frac{\partial f(a)}{\partial a}$ *is positive iff* $V_1^c(a, f(a)) - V_1^s(a, f(a))$ *and* $V_2^s(a, f(a)) - V_2^c(a, f(a))$ *have the same sign.*

Lemma 2 is self-evident from equation (7). It essentially tells us that the relationship between areas that are optimal for CT and ST is driven by the rate of change of value functions at the crossing points.

However, for a model to be in line with the data evidence presented in the previous section, not only should the slope be positive $\frac{\partial f(a)}{\partial a} > 0$ but the differences in the nominator and denominator should be positive themselves. Namely, the first derivative of the value function with respect to assets should be higher for CT than for ST, while the first derivative of the value function with respect to income should be lower for CT than for ST. If so, the value function of CT increases faster with assets at the crossing point and thus it will be higher than the value function of ST. The reverse holds for income. This argument is summarized in Proposition 2 (the proof directly follows from the fact that $f(a)$ is unique).

Proposition 2: *If* $V_1^c(a, f(a)) > V_1^s(a, f(a))$ *and* $V_2^s(a, f(a)) > V_2^c(a, f(a))$*, then for all pairs* $(a, f(a))$ *the following holds*
$\forall y \geq f(a)\ V^s(a, y) > V^c(a, y)$ *and* $\forall y < f(a)\ V^s(a, y) < V^c(a, y)$,
$\forall a \geq f^{-1}(y)\ V^s(a, y) < V^c(a, y)$ *and* $\forall a < f^{-1}(y)\ V^s(a, y) > V^c(a, y)$.

It is easy to draw a connection between Proposition 1 and Proposition 2. Proposition 1 essentially claims that the slope $\frac{\partial f(a)}{\partial a}$ is positive and equals a constant. In fact, if one takes $V^c(a, y) = W^c(a, y)$ and $V^s(a, y) = W^s(a, y)$ from the toy model, the condition $V^c(a, f(a)) = V^s(a, f(a))$ reduces to the condition $\kappa \frac{\partial \lambda}{\partial a} = 0$, where λ is a Lagrange multiplier on the non-negativity constraint of the consumption targeter $c_1 + c_2 > a + 2\underline{y} + \mu$ and $\kappa > 0$. Consequently, $\frac{\partial f(a)}{\partial a} = \frac{1}{(2\beta^2 + 2\beta - 1)} > 0$ for $\beta > \frac{\sqrt{3}-1}{2} \approx 0.37$ (for the derivations see Appendix C).

Stable choices and Switchers

From now on, I study the consumption response to an exogenous policy intervention that provides a positive change $\Delta > 0$ in assets a for all agents in the economy. This policy intervention could be thought of as a large-scale tax rebate program that took place several times in the US. I assume that all agents in the economy make the choice (6) at least two times, before and after the change; and the slope $\frac{\partial f(a)}{\partial a}$ is positive. Then, based on Proposition 2, there are three possible responses in the type choice to the change in assets. There could be *stable consumption targeters* for whom

$$V^c(a, y) > V^s(a, y) \quad \textit{and} \quad V^c(a + \Delta, y) > V^s(a + \Delta, y),$$

stable savings targeters for whom

$$V^c(a, y) < V^s(a, y) \quad \textit{and} \quad V^c(a + \Delta, y) < V^s(a + \Delta, y),$$

and *switchers* for whom

$$V^c(a, y) < V^s(a, y) \quad \textit{and} \quad V^c(a + \Delta, y) > V^s(a + \Delta, y).$$

Marginal Propensity to Consume

Imagine that a policymaker is interested in estimating the marginal propensity to consume (MPC) out of asset change Δ, which is the ratio between consumption change due to the policy intervention and the asset change Δ. Based on the three possible responses in the type choice to the change in assets, there are three types of MPC measures. The MPC of a *stable consumption targeter* is

$$\frac{c(a + \Delta, y) - c(a, y)}{\Delta},$$

the MPC of a *stable savings targeter* is

$$\frac{[y - s(a + \Delta, y)] - [y - s(a, y)]}{\Delta},$$

the MPC of *switchers* is

$$\frac{c(a + \Delta, y) - [y - s(a, y)]}{\Delta} = \frac{c(a + \Delta, y) - c(a, y)}{\Delta} + \frac{c(a, y) + s(a, y) - y}{\Delta}.$$

In the textbook consumption-savings models, the sum $c(a, y) + s(a, y)$ coincides with current income y, so the MPC of switchers coincides with one of the consumption targeters. But if consumption and savings are chosen under different optimization problems, this may

be not the case. For example, in the toy model presented above, constrained consumption plans and savings plans do not sum up to an income state. This also happens in the toy model solved for CARA utility, because the utility function exhibits prudence, and thus policy functions are responsive to uncertainty, which is treated differently by CT and ST. The difference $c(a,y)+s(a,y)-y$ in this case is equal to $\frac{(R-R^2)\sigma^2}{2(1+R+R^2)}$, where R is an interest rate on assets and σ^2 is income variance. The difference $(R-R^2)$ appears exactly because ST splits uncertainty while CT bundles it. For the full derivations see Appendix D.

Average Marginal Propensity to Consume

Average MPC is defined as $m(\Delta)=\int\int m(\Delta;a,y)f(a+\Delta,y)dady$ (where $f(a,y)$ is a density function) and contains MPCs of consumption targeters, savings targeters, and switchers. For the economy of consumption targeters and savings targeters and $\Delta>0$ the average MPC takes the form of

$$\begin{aligned} m(\Delta) = & \underbrace{\int\int_{(a,y)\&(a+\Delta,y)\in S^c} \frac{c(a+\Delta,y)-c(a,y)}{\Delta} f(a+\Delta,y)dady}_{\text{consumption targeters}} \\ & - \underbrace{\int\int_{(a,y)\&(a+\Delta,y)\in S^s} \frac{s(a+\Delta,y)-s(a,y)}{\Delta} f(a+\Delta,y)dady}_{\text{savings targeters}} \\ & + \underbrace{\int\int_{(a,y)\in S^s\&(a+\Delta,y)\in S^c} \frac{c(a,y)+s(a,y)-y}{\Delta} f(a+\Delta,y)dady}_{\text{Wedge}} \end{aligned} \tag{8}$$

where S^c is the area in the state space where the agent is a consumption targeter, S^s is the area in the state space where the agent is a savings targeter, $s(a,y)$ is the policy function of savings targeters, $f(a,y)$ is the equilibrium distribution of agents in the state space.

I call the *wedge* the following expression

$$\int\int_{(a,y)\in S^s\&(a+\Delta,y)\in S^c} \frac{c(a,y)+s(a,y)-y}{\Delta} f(a+\Delta,y)dady.$$

The first two expressions in equation (8) define an average slope of policy functions in the economy. To the best of my knowledge, any research on average marginal propensity to consume models it as the average slope of policy function, disregarding the last component.

To summarize, the empirically observed average MPC from an economy-wide transfer program (e.g. Tax Rebate Program of 2008 in the US) $m(\Delta)$ is, in theory, composed of the average slope of policy function (first two components in equation (8)) and the wedge (last component in equation (8)). The wedge affects not only the size of the average slope

of the policy function required to match empirically observed MPC but also the properties of policy functions. For example, if the wedge is concave in assets and negative, one will need a more concave slope of policy functions to match the concavity of observed MPC without accounting for the wedge. In other words, without accounting for the wedge, one will overestimate or underestimate the concavity of the policy functions depending on the properties of the wedge.

Even though, in Section 3, I study the wedge in the average MPC induced by switchers using a specific model of the choice between consumption and savings plans, the formula for average MPC is general. I argue that the wedge will be generated in any model where consumption and savings plans do not provide the same utility, agents can choose between plans, and $s + c \neq y$. This argument is summarized in Proposition 3 (the proof directly follows from equation (8)).

Proposition 3: *If* $\exists(a, y)$ *such that* $c(a, y) + s(a, y) - y \neq 0$ *and there are agents for whom* $V^c(a, y) < V^s(a, y) \quad \& \quad V^c(a + \Delta, y) > V^s(a + \Delta, y) \quad or$
$V^c(a, y) > V^s(a, y) \quad \& \quad V^c(a + \Delta, y) < V^s(a + \Delta, y)$*, then*
there is a wedge between the average MPC and the average slope of policy functions in economy.

3 Quantitative model and wedge assessment

In this section, I extend the toy model from Section 1 to be a fully-dynamical infinite-horizon model. The aim of the fully-dynamical infinite-horizon model is to assess the wedge induced by the switchers in the economy.

Infinite-horizon model with infrequent updating

If the agent updates her plans infrequently, consumption and savings decision variables provide different utility flows. The agent chooses between consumption and savings plans by comparing maximized expected lifetime utilities. There is a continuum of infinitely-lived agents endowed with assets a and income y. I assume that the agent updates her plans each d periods, and thus they create plans for d periods ahead based on *today's* information and revise the plan in d periods.

$$\mathcal{V}^c(a, y) = max_c \sum_{t=1}^{d} \beta^{t-1} u(c_t) + \beta^d E_y[\mathcal{V}(a', y')] \quad s.t. \tag{9}$$

$$a' = a + \sum_{t=1}^{d} (y_t - c_t),\ c_t > 0,\ a' > 0\ \forall t$$

The savings plan is determined by the following optimization problem:

$$\mathcal{V}^s(a,y) = max_s E_y[\sum_{t=1}^{d} \beta^{t-1} u(y_t - s_t)] + \beta^d \mathcal{V}(a',y')] \quad s.t. \tag{10}$$

$$a' = a + \sum_{t=1}^{d} s_t,\ y_t - s_t > 0,\ a' > 0\ \forall t$$

The agent chooses between consumption and savings plan by comparing expected utilities from the two types of plans:

$$\Big[\mathcal{V}(a,y), I(a,y)\Big] = max\{\mathcal{V}^s(a,y), \mathcal{V}^c(a,y)\},$$

where $I(a,y)$ is type-decision rule.

The consumption policy function is a set of d policy functions that determine consumption in each first d periods $c(a,y) = [c_1(a,y)...c_d(a,y)]$. The savings policy function is a set of d policy functions that determine savings in each first d periods $s(a,y) = [s_1(a,y)...s_d(a,y)]$, where $s \equiv y - c$.

The first-order conditions for CT problem for all $t \leq d$ are

$$u'(c_t) = \beta^{d-t} E_0[\mathcal{V}_1(a',y')] \tag{11}$$

Essentially, the consumer follows a deterministic path for consumption in the first d periods and calculates it by scaling consumption in the first period. For the CRRA utility function, $c_{t+1} = \beta^{-\gamma} c_t$. To keep consumption certain, a consumption targeter bears high uncertainty in the continuation value $\mathcal{V}(a',y')$, which has both components a' and y' uncertain for the consumer.

The first-order conditions for ST problem for all $t \leq d$ are

$$E_0[u'(y_t - s_t)] = \beta^{d-t} E_0[\mathcal{V}_1(a',y')] \tag{12}$$

While making a decision regarding s_t for $t > 1$, a savings targeter accounts for uncertainty in y_t. By doing so, she relaxes the uncertainty in the continuation value $\mathcal{V}(a',y')$, which has an uncertain income y' but certain assets a' because the saver has explicitly chosen their level.

The story behind the model could be the following. The agent receives paychecks *d times* a month to her bank account and goes shopping on paydays. Meanwhile, she plans her spending only *once* a month and does not violate the plan. Because the agent goes shopping *d times* but chooses spending amounts only *once*, she should create an advance plan for all of the shopping trips before knowing her paychecks. Because future paychecks are uncertain

and the agent cannot violate her plan, she decides how to absorb unpredictable changes in the paychecks. She might do it with savings, so she chooses the amounts to spend in a shop, and thus, saves less (more) if her paycheck is lower (higher) than expected. On the other hand, she might do it with consumption, so she chooses the amount to save for the next month, and thus, spends less (more) in a shop if her paycheck is lower (higher) than expected. The main constraint the agent faces is that even for the realization of d minimal possible paychecks, she has to consume something and hold assets for the next month.

It is worth mentioning that a savings targeter learns about her income at the moment of its realization because consumption is defined as $c \equiv y - s$ and she has chosen the amount of savings. This fact does not put savings targeters in a privileged position because both types commit to a plan and the saver learns her income only when it is irrelevant to her decisions.

The results have clear intuition: CT secures against the risk in the third period (which is more sensitive to her in comparison to ST) if she has a lot of assets, while ST secures against the risk in the second period (more sensitive to her) if she has a high-income state.

A stationary equilibrium for this economy consists of individual policy functions $c(a, y) = [c_1(a, y)...c_d(a, y)]$, $s(a, y) = [s_1(a, y)...s_d(a, y)]$, type-decision rule $I(a, y)$ and stationary probability distribution ψ over assets with transition matrix $P(a, y_t...y_{t+d})$. Everyone updates plans identically infrequently. Shares of consumption targeters and savings targeters are defined in equilibrium based on ψ and type-decision rule $I(a, y)$.

Income process and benchmark calibration

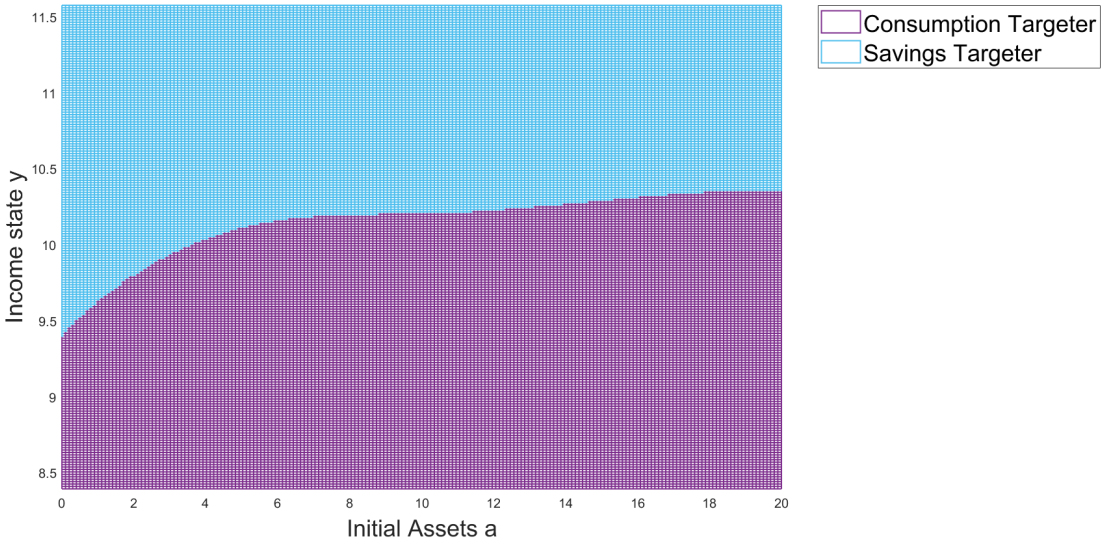
Income risk is the crucial source of heterogeneity that drives the difference between CT and ST types. By assuming an earnings process with time-invariant persistence and normally distributed shocks, I would put too many modeling restrictions on income heterogeneity. Instead, I use a non-parametric approach to approximate nonlinear, non-normal, and time-dependent stochastic processes as a first-order Markov chain following Arellano et al. (2017) and De Nardi et al. (2019). I approximate the process on the SCE data that was described in Section 1.

In this quantitative exercise, I assume CRRA utility funtion $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$ with $\gamma = 3$ as a benchmark. Time discounting parameter β is set equal to 0.97. I calibrate the model at the quarterly frequency and set a benchmark updating frequency $d = 4$, meaning that the agents do not update plans for one year. I provide extensive comparative statics for different parameter values and show that all the results are stable.

Optimal plans

Figure 3 depicts the optimal choice of types for different income states y and initial assets a. One can see that the areas of CT and ST are split by a curve with a positive slope, so for given assets a STs are those with a *higher* income state y but for a given income state y STs are those with *lower* assets a. Further, the curvature of the line is mainly concentrated in the low asset levels.

Figure 3: Optimal Plans for infinite-horizon model

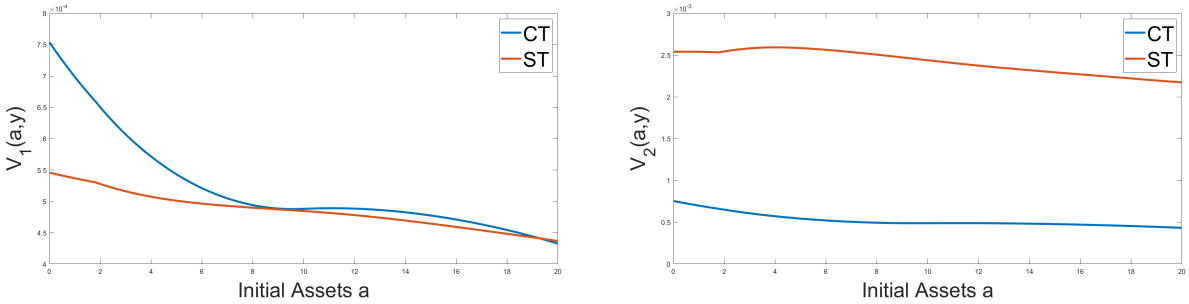


Intuitively, the decision between consumption and savings plans is driven by the trade-off between uncertainty in assets and uncertainty in income. Because a savings targeter uses savings as a decision variable, she knows with certainty the level of assets in the period of update $a' = a + \sum_{t=1}^{d} s_t$. Thus, for the saver, income realization in the second period (period of inattentiveness) does not affect accumulated assets and is fully absorbed by consumption. Alternatively, if an agent uses consumption as a decision variable, she knows with certainty her consumption level in each period and income realization in the second period is absorbed by her assets $a' = a + \sum_{t=1}^{d}(y_t - c_t)$. Thus, an agent with high (low) assets tends to be a consumption targeter (saver).

The results could be better described through the lens of Proposition 2, which relates the type choice to the first derivatives of value functions with respect to income and assets. In Figure 4, one can see that CT has a higher first derivative of value function with respect to income than ST and the discrepancy is mainly concentrated in the low asset levels. In contrast, ST has a higher first derivative of value function with respect to assets than CT and the discrepancy is stable. CT values additional assets more than ST because they serve

as the main source of security against uncertainty in consumption for her, and the source is the most valuable when the initial assets are low.

On the other hand, ST experiences uncertainty in her consumption during d periods of inattentiveness, and the only source of security for her is a high-income state. ST values a higher-income state more than CT, who also uses assets to secure against consumption uncertainty because, given the high-income state today and some persistence in the income process, she is less likely to experience a low-income state during d periods of inattentiveness.



(a) 1a (b) 1b

Figure 4

Table 3 contains the slope of $f(a)$ over assets distribution for different parameter values. The curvature of the line is concentrated in the first quartile of distribution for any parameter changes but is more or less pronounced depending on them. Firstly, the curvature increases when time discounting β is lower, so the agents are less patient. This happens because the difference between postponing risk (CT) and not (ST) is more pronounced when β is lower. Secondly, the slope increases when risk-aversion γ is lower because in this case the difference between bundling risk (CT) and splitting it (ST) becomes more pronounced. Thirdly, the slope decreases when the agents update their plans more frequently (the types are the same when $d = 1$). Lastly, the curvature is lower for a more persistent income process, because the overall uncertainty the agent faces is lower in this case.

Table 3: Slope of $f(a)$ over assets distribution

Assets distibution	Quartile 1	Quartile 2	Quartile 3	Quartile 4
Benchmark model	0.037	0.008	0.006	0.002
Less patient: $\beta = 0.94$	0.183	0.062	-0.000	0.009
Less risk-averse: $\gamma = 1.5$	0.080	0.029	0.002	0.005
More frequent updating: $d = 2$	0.019	0.003	0.004	0.002
More persistent income (CEX data)	0.024	0.001	0.003	-0.000

Marginal Propensity to Consume and wedge assessment

In this step, I aim to calibrate the wedge in average MPC induced by the switchers in order to assess its size. I calculate the average MPC out of assets in the economy of consumption targeters and savings targeters based on the formula (8). Table 4 contains the main result of the paper. It shows that MPC is increasing in assets when switchers are present and decreasing when they are excluded.

Table 4: Marginal propensity to consume (MPC) over assets distribution

Assets distibution	Quartile 1	Quartile 2	Quartile 3	Quartile 4
Benchmark model				
Average MPC	0.007	0.025	0.028	0.023
Average MPC (without switchers)	0.089	0.031	0.029	0.023

Description of table 4.

The results in Table 5 demonstrate that the wedge induced by switchers is non-negligible and is 20% of the average MPC under the benchmark calibration.

Table 5: Ratio of **wedge** to average slope of policy functions for $d = 2$

Model	Benchmark	$\beta = 0.94$	$\gamma = 1.5$	CEX data
$c(a, y) + s(a, y) - y$ (negative)	123.3	192.9	228.4	106.5
Avegare slope	0.066	0.095	0.100	0.067
Wedge (negative)	0.013	0.004	0.005	0.006
Ratio	0.201	0.043	0.055	0.086

Note: Benchmark: $\gamma = 3, \beta = 0.97$

In the first row of Table 5, I show the absolute values of the discrepancy between $c(a, y) + s(a, y)$ and y, which are higher for lower β and γ and lower for slightly more persistent income (CEX data). This happens because lower β and γ lead to more curved $f(a)$ and the policy functions relate to first derivatives of value functions through FOCs and Envelope Theorem ($c(a, y)$ relates to $V_1^c(a, y)$ through FOC (11) and Envelope Theorem, $s(a, y)$ relates to $V_1^s(a, y)$ through FOC (12) and Envelope Theorem). Average slope changes naturally with parameter changes because for lower β and γ current consumption is more valuable because an agent is less patient (lower β) or the precautionary motive is less pronounced (lower γ). I leave a more formal analysis for future research.

Even though the absolute values of the discrepancy $c(a, y) + s(a, y) - y$ are higher for

lower β and γ, it is not the case for the wedge itself. This happens because the wedge

$$\int\int_{(a,y)\in S^s \& (a+\Delta,y)\in S^c} \frac{c(a,y)+s(a,y)-y}{\Delta} f(a+\Delta,y)dady.$$

also contains the density measure of agents who switched between the types, which depends on the levels of $f(a)$. Thus, for proper calibration of the wedge, both the size of the discrepancy $c(a,y)+s(a,y)-y$ and density function matters.

One can see that for any parameter values the wedge amounts to a sizeable share of the average MPC. Thus, matching the average MPC and the slope of policy functions is fallacious. As the wedge is negative, failing to account for it leads to the underestimation of the required slope of policy functions required to match the observed average MPC.

Importantly, the wedge is concentrated in the first quartile of asset distribution because of the relation between the discrepancies $c(a,y)+s(a,y)-y$ and $V_1^c(a,y)-V_1^s(a,y)$ (Figure 4) through FOCs and Envelope Theorem. Thus, the mismatch will be the highest for the poorest agents who typically have high MPC and are crucial for redistribution policy.

Conclusion

This paper presents a novel explanation for the conflicting empirical findings on the relationship between the marginal propensity to consume (MPC) and liquid assets. While standard consumption-savings models predict a concave consumption policy function, implying a decreasing MPC with increasing liquid assets, recent empirical evidence (Lewis et al., 2019) suggests a positive correlation. By incorporating decision-variable switching and infrequent financial plan updates into a standard income fluctuation model, this paper reconciles these empirical findings with theoretical predictions without deviating from the fundamental concavity of individual consumption functions.

The key mechanism driving this result is the existence of two distinct agent types: consumption targeters, who optimize consumption directly, and savings targeters, who optimize savings decisions. The model predicts that low-asset, high-income agents are more likely to adopt a savings-targeting strategy, while high-asset, low-income agents tend to become consumption targeters. Importantly, some agents switch between these two decision rules in response to small changes in liquid assets. This switching behavior introduces an additional, negative component to the aggregate MPC, which is more pronounced at low levels of liquid assets. Consequently, the empirical relationship between aggregate MPC and liquid assets can be positive if the negative switching component is sufficiently large. A calibrated infinite-horizon model demonstrates that this switching-induced component is nontrivial, accounting for approximately 20% of the observed MPC under the benchmark specification.

The results have significant policy implications. In particular, fiscal stimulus policies that assume a uniform MPC across households may overestimate the effectiveness of liquidity injections for low-asset households, as a portion of these households may respond by switching from a savings-targeting to a consumption-targeting strategy rather than increasing immediate consumption. Future research could extend this framework to incorporate life-cycle considerations, exploring how the propensity to switch between decision rules evolves with age and wealth accumulation.

By bridging the gap between standard theory and empirical observations, this paper provides a deeper understanding of the determinants of MPC and contributes to the broader literature on consumption behavior, inattention, and liquidity constraints. The findings underscore the necessity of accounting for heterogeneity in decision-making processes when analyzing aggregate consumption dynamics and designing economic policy interventions.

References

Andre, P., G. Nikolakoudis, K. Sastry, and J. Flynn (2024). Quick-fixing: Near-rationality in consumption and savings behavior. Working paper.

Arellano, M., R. Blundell, and S. Bonhomme (2017). Earnings and consumption dynamics: A nonlinear panel data framework. *Econometrica 85*(3), 693–734.

Bartoš, V., M. Bauer, J. Chytilová, and F. Matějka (2016). Attention discrimination: Theory and field experiments with monitoring information acquisition. *American Economic Review 106*(6), 1437–75.

Bordalo, P., N. Gennaioli, and A. Shleifer (2013). Salience and asset prices. *American Economic Review 103*(3), 623–28.

Boutros, M. (2022). Windfall Income Shocks with Finite Planning Horizons. Staff Working Papers 22-40, Bank of Canada.

Broda, C. and J. A. Parker (2014). The economic stimulus payments of 2008 and the aggregate demand for consumption. *Journal of Monetary Economics 68*, S20–S36.

Carroll, C. D. and M. S. Kimball (1996). On the concavity of the consumption function. *Econometrica 64*(4), 981–992.

Carroll, C. D. and M. S. Kimball (2001, October). Liquidity constraints and precautionary saving. Working Paper 8496, National Bureau of Economic Research.

Chetty, R., A. Looney, and K. Kroft (2009). Salience and taxation: Theory and evidence. *American Economic Review 99*(4), 1145–77.

De Nardi, M., G. Fella, and G. Paz-Pardo (2019). Nonlinear Household Earnings Dynamics, Self-Insurance, and Welfare. *Journal of the European Economic Association 18*(2), 890–926.

Deaton, A. (1991). Saving and liquidity constraints. *Econometrica 59*(5), 1221–1248.

Fagereng, A., M. B. Holm, and G. J. Natvik (2021). Mpc heterogeneity and household balance sheets. *American Economic Journal: Macroeconomics 13*(4), 1–54.

Gabaix, X. (2014). A sparsity-based model of bounded rationality. *The Quarterly Journal of Economics 129*(4), 1661–1710.

Gabaix, X. (2016). Behavioral macroeconomics via sparse dynamic programming. NBER Working Papers 21848, National Bureau of Economic Research, Inc.

Gabaix, X. and D. Laibson (2002). The 6d bias and the equity premium puzzle. Harvard Institute of Economic Research Working Papers 1947, Harvard - Institute of Economic Research.

Jappelli, T. and L. Pistaferri (2014). Fiscal policy and mpc heterogeneity. *American Economic Journal: Macroeconomics 6*(4), 107–36.

Kaplan, G. and G. L. Violante (2022). The marginal propensity to consume in heterogeneous agent models. Working Paper 30013, National Bureau of Economic Research.

Kimball, M. S. (1990). Precautionary saving in the small and in the large. *Econometrica 58*(1), 53–73.

Kőszegi, B. and A. Szeidl (2013). A model of focusing in economic choice. *The Quarterly journal of economics 128*(1), 53–104.

Kueng, L. (2018). Excess sensitivity of high-income consumers. *Quarterly Journal of Economics 133*, 1693–1751.

Kőszegi, B. and F. Matějka (2020). Choice simplification: A theory of mental budgeting and naive diversification. *The Quarterly Journal of Economics 135*(2), 1153–1207.

Lewis, D. J., D. Melcangi, and L. Pilossoph (2019). Latent heterogeneity in the marginal propensity to consume. *FRB of New York Staff Report* (902).

Miranda-Pinto, J., D. P. Murphy, K. Walsh, and E. Young (2020, February). Saving Constraints, Debt, and the Credit Market Response to Fiscal Stimulus. Working Papers 20-07, Federal Reserve Bank of Cleveland.

Mondria, J. and C. Quintana-Domeque (2013). Financial contagion and attention allocation. *The Economic Journal 123*(568), 429–454.

Pope, D. G. (2009). Reacting to rankings: Evidence from "America's Best Hospitals". *Journal of Health Economics 28*(6), 1154–1165.

Reis, R. (2006). Inattentive consumers. *Journal of Monetary Economics 53*(8), 1761–1800.

Shapiro, M. D. and J. Slemrod (2009). Did the 2008 tax rebates stimulate spending? *American Economic Review 99*(2), 374–79.

Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics 50*(3), 665 – 690.

Tutino, A. (2013). Rationally inattentive consumption choices. *Review of Economic Dynamics 16*(3), 421–439.

Zeldes, S. P. (1989). Consumption and liquidity constraints: An empirical investigation. *Journal of Political Economy 97*(2), 305–346.

Zhang, C. Y., A. B. Sussman, N. Wang-Ly, and J. K. Lyu (2022). How consumers budget. *Journal of Economic Behavior & Organization 204*, 69–88.

Appendix A

Proof of Lemma 1.
The solution to unconstrained CT problem is the following:

$$c_1^* = \frac{\beta^2(a+3\mu)}{1+\beta+\beta^2}$$

$$c_2^* = \frac{\beta(a+3\mu)}{1+\beta+\beta^2}$$

First-period utility is $u(c_1^*) = -(\frac{\beta^2(a+3\mu)}{1+\beta+\beta^2})^2$
Second-period utility: $u(c_2^*) = -\beta(\frac{\beta(a+3\mu)}{1+\beta+\beta^2})^2$
Third-period utility
$E[u(a+y_3+y_1+y_2-c_1^*-c_2^*)] = -\beta^2(2\sigma^2+4\mu^2+4\mu(\frac{\beta(a+(1-2\beta-2\beta^2)\mu)}{1+\beta+\beta^2})+((\frac{\beta(a+(1-2\beta-2\beta^2)\mu)}{1+\beta+\beta^2}))^2)$.

Solution to unconstrained ST problem is the following:

$$s_1^* = \frac{-\beta^2 a + (1+\beta-2\beta^2)\mu)}{1+\beta+\beta^2}$$

$$s_2^* = \frac{-\beta a + (\beta-1)^2\mu)}{1+\beta+\beta^2}$$

First-period utility: $u(\mu - s_1^*) = (\mu - \frac{-\beta^2 a+(1+\beta-2\beta^2)\mu)}{1+\beta+\beta^2})^2$
Second-period utility: $E[u(y_2 - s_2^*)] = -\beta(\sigma^2+\mu^2-2\mu\frac{-\beta^2 a+(1+\beta-2\beta^2)\mu)}{1+\beta+\beta^2}+(\frac{-\beta^2 a+(1+\beta-2\beta^2)\mu)}{1+\beta+\beta^2})^2)$
Third-period utility: $E[u(a+y_3+s_1^*+s_2^*)] = -\beta^2(\sigma^2+\mu^2+2\mu\frac{a+(2-\beta-\beta^2)\mu}{1+\beta+\beta^2}+(\frac{a+(2-\beta-\beta^2)\mu}{1+\beta+\beta^2})^2)$

Simplifying the expressions one will get $W^c - W^s = (\beta-\beta^2)\sigma^2$.

Appendix B

Proof of Proposition 1.
The logic of the proof is the following. I measure the difference between constrained and unconstrained V^c at the point $\hat{a}-\xi$, where $\hat{a}$ is the point where the constraint binds ($V^c_{un}(\hat{a}-\xi)$ is not feasible though). I show that the difference is increasing in ξ and does not depend on a or μ. Then, I do exactly the same for $\hat{\mu}+\omega$, where $\hat{\mu}$ is the point where the constraint binds. For large enough ξ and ω the difference between constrained and unconstrained V^c becomes larger than $(\beta-\beta^2)\sigma^2$, so ST becomes preferable. As the difference does not depend on a or μ, the function that separates the optimal type choices is a line.

Optimal solutions:

$c_1 = c_1^* - \frac{(1+\beta-\beta^2)\lambda}{1+\beta+\beta^2}$

$c_2 = c_2^* - \frac{(1-\beta+\beta^2)\lambda}{1+\beta+\beta^2}$

$\lambda = \frac{a+(1-2\beta-2\beta^2)\mu+2\underline{y}(1+\beta+\beta^2)}{2}$

The point at which the non-negativity constraint binds

$$\hat{a} = -2\underline{y}(1+\beta+\beta^2) + \hat{\mu}(2\beta+2\beta^2-1)$$

$$\hat{\mu} = \frac{\hat{a}+2\underline{y}(1+\beta+\beta^2)}{2\beta+2\beta^2-1}$$

1) Take $\underline{a} = \hat{a} - \xi$

$c_1(\hat{a}-\xi) = c_1^*(\hat{a}-\xi) + \frac{(1+\beta-\beta^2)\xi}{2(1+\beta+\beta^2)}$

$c_2(\hat{a}-\xi) = c_2^*(\hat{a}-\xi) + \frac{(1-\beta+\beta^2)\xi}{2(1+\beta+\beta^2)}$

$a+\mu-c_1-c_2 = a+\mu-c_1^*-c_2^* - \frac{\xi}{1+\beta+\beta^2}$

$c_1^*(\hat{a}-\xi) = 2\beta^2(\hat{\mu}-\underline{y}) - \frac{\beta^2\xi}{1+\beta+\beta^2}$

$c_2^*(\hat{a}-\xi) = 2\beta(\hat{\mu}-\underline{y}) - \frac{\beta\xi}{1+\beta+\beta^2}$

$a+\mu-c_1^*-c_2^* = -\underline{y} + \frac{\beta\xi}{1+\beta+\beta^2} + \frac{\beta^2\xi}{1+\beta+\beta^2}$

$U_1(c_1^*) - U_1(c_1) = \frac{-(1+\beta-\beta^2)\xi}{2(1+\beta+\beta^2)}(2c_1^* + \frac{(1+\beta-\beta^2)\xi}{2(1+\beta+\beta^2)})$

$U_2(c_1^*) - U_2(c_1) = \frac{-(1-\beta+\beta^2)\xi}{2(1+\beta+\beta^2)}(2c_2^* + \frac{(1-\beta+\beta^2)\xi}{2(1+\beta+\beta^2)})$

$U_3(c_3^*) - U_3(c_3) = 4\mu\frac{\xi}{1+\beta+\beta^2} + \frac{\xi}{1+\beta+\beta^2}(2(\mu+a-c_1^*-c_2^*) + \frac{\xi}{1+\beta+\beta^2})$

$U_1(c_1^*) - U_1(c_1) + \beta(U_2(c_1^*) - U_2(c_1)) + \beta^2(U_3(c_3^*) - U_3(c_3))$ does not depend on $\hat{\mu}$

2) The same procedure for μ, take $\underline{\mu} = \hat{\mu} + \omega$.

$c_1^*(\hat{\mu}) = \frac{2\beta^2(\hat{a}+3\underline{y})}{2\beta^2+2\beta-1}$

$c_2^*(\hat{\mu}) = \frac{2\beta(\hat{a}+3\underline{y})}{2\beta^2+2\beta-1}$

$c_1(\hat{\mu}) = c_1^*(\hat{\mu}) + \frac{(1+\beta-\beta^2)\omega}{2(1+\beta+\beta^2)}$

$c_2(\hat{\mu}) = c_2^*(\hat{\mu}) + \frac{(1-\beta+\beta^2)\omega}{2(1+\beta+\beta^2)}$

$c_1^*(\hat{\mu}+\omega) = c_1^*(\hat{\mu}) + \frac{3\beta^2\omega}{1+\beta+\beta^2}$

$c_2^*(\hat{\mu}+\omega) = c_2^*(\hat{\mu}) + \frac{3\beta\omega}{1+\beta+\beta^2}$

$a+\mu-c_1-c_2 = a+\mu-c_1^*-c_2^* - \frac{\omega}{1+\beta+\beta^2}$

$a+\mu-c_1^*-c_2^* = -2\underline{y} - \frac{3\omega(\beta+\beta^2)}{1+\beta+\beta^2}$

$U_1(c_1^*) - U_1(c_1) = -\frac{\omega(1+\beta-\beta^2)}{2(1+\beta+\beta^2)}(2c_1^* + \frac{(1+\beta-\beta^2)\omega}{2(1+\beta+\beta^2)})$

$U_2(c_2^*) - U_2(c_2) = -\frac{\omega(1-\beta+\beta^2)}{2(1+\beta+\beta^2)}(2c_2^* + \frac{(1-\beta+\beta^2)\omega}{2(1+\beta+\beta^2)})$

$U_3(c_3^*) - U_3(c_3) = 4\frac{\hat{a}+2\underline{y}(1+\beta+\beta^2)}{2\beta+2\beta^2-1}\frac{\omega}{1+\beta+\beta^2} + \frac{\omega}{1+\beta+\beta^2}(2(\mu+a-c_1^*-c_2^*) + \frac{\omega}{1+\beta+\beta^2})$

$U_1(c_1^*) - U_1(c_1) + \beta(U_2(c_1^*) - U_2(c_1)) + \beta^2(U_3(c_3^*) - U_3(c_3))$ does not depend on $\hat{a}$

Appendix C

The agent might create a consumption plan (use consumption as a decision variable):

$$W^c(a, y_1) = \max_c u(c_1) + \beta u(c_2) + \beta^2 E[u(a + y_1 + y_2 + y_3 - c_1 - c_2)] \quad s.t. \tag{13}$$

$$c_1 > 0,\ c_2 > 0,\ c_1 + c_2 \leq a + 2\underline{y} + y_1$$

The agent might use savings as a decision variable:

$$W^s(a, y_1) = \max_s u(y_1 - s_1) + E[\beta u(y_2 - s_2) + \beta^2 u(a + y_3 + s_1 + s_2)] \quad s.t. \tag{14}$$

$$s_1 \leq y_1,\ s_2 \leq \underline{y},\ a + \underline{y} + s_1 + s_2 > 0$$

$f(a)$ is a function such that

$$W^c(a, f(a)) = W^s(a, f(a))$$

Take the full derivative

$$\frac{dW^c(a, f(a))}{da} = \frac{dW^s(a, f(a))}{da} \tag{15}$$

For ST constraints never bind, so by Envelope Theorem a has only direct effect on $W^s(a, f(a))$, so $\frac{dW^s(a,f(a))}{da} = \frac{\partial c_1^*}{\partial a}\frac{\partial f(a)}{\partial a} + \beta E[\frac{\partial c_2^*}{\partial a}\frac{\partial f(a)}{\partial a}] + \beta^2 E[\frac{\partial c_3^*}{\partial a}(1 + \frac{\partial f(a)}{\partial a})]$
Solutions to CT's problem:

$$c_1 = \frac{\beta^2(a + 3\mu)}{1 + \beta + \beta^2} - \frac{1 - \beta + \beta^2}{1 + \beta + \beta^2}\lambda = c_1^* - k_1\lambda$$

$$c_2 = \frac{\beta(a + 3\mu)}{1 + \beta + \beta^2} - \frac{1 + \beta - \beta^2}{1 + \beta + \beta^2}\lambda = c_2^* - k_2\lambda$$

$$\lambda = \frac{a + (1 - 2\beta - 2\beta^2)\mu + 2\underline{y}(1 + \beta + \beta^2)}{2}$$

$\frac{dW^c(a,f(a))}{da} = [\frac{\partial u}{\partial c_1^*} - k_1\lambda]\frac{\partial c_1^*}{\partial a} + \beta[\frac{\partial u}{\partial c_2^*} - k_2\lambda]\frac{\partial c_2^*}{\partial a} - [\frac{\partial u}{\partial c_1^*} - k_1\lambda]\frac{\partial \lambda}{\partial a} - \beta[\frac{\partial u}{\partial c_2^*} - k_2\lambda]\frac{\partial \lambda}{\partial a} + \beta^2[\frac{\partial u}{\partial c_3^*} + (k_1+k_2)\lambda][-\frac{\partial c_1^*}{\partial a} - \frac{\partial c_2^*}{\partial a} + (k_1+k_2)\frac{\partial \lambda}{\partial a} + 1 + 3\frac{\partial f(a)}{a}$
$\frac{\partial c_1^*}{\partial a}[\frac{\partial u}{\partial c_1^*} - k_1\lambda - \beta^2\frac{\partial u}{\partial c_3^*} - \beta^2(k_1+k_2)\lambda]$

From FOC $\frac{\partial u}{\partial c_1^*} = \beta^2\frac{\partial u}{\partial c_3^*}$. $\frac{\partial c_1^*}{\partial a} = \frac{\beta^2(1+3\frac{\partial f(a)}{a})}{1+\beta+\beta^2}$. $-k_1\lambda - (k_1+k_2)\lambda = -\lambda$
$\frac{\partial c_2^*}{\partial a}[\beta\frac{\partial u}{\partial c_2^*} - k_2\beta\lambda - \beta^2\frac{\partial u}{\partial c_3^*} - \beta^2(k_1+k_2)\lambda]$

From FOC $\beta\frac{\partial u}{\partial c_2^*} = \beta^2\frac{\partial u}{\partial c_3^*}$. $\frac{\partial c_1^*}{\partial a} = \frac{\beta(1+3\frac{\partial f(a)}{\partial a})}{1+\beta+\beta^2}$. $-k_2\beta\lambda - \beta^2(k_1+k_2)\lambda = -\beta\lambda$
$\frac{\partial \lambda}{\partial a} = \frac{1+(1-2\beta-2\beta^2)\frac{\partial f(a)}{\partial a}}{1+\beta+\beta^2}$

Because $c_1^* = \mu - s_1^*$, using FOCs of ST problem and Envelope Theorem one could show that $\frac{dW^s(a,f(a))}{da} = \beta^2\frac{\partial u}{\partial c_3^*}[1+3\frac{\partial f(a)}{a}]$, which means that direct effects of unconstrained CT and ST problems coincide.
$k_1 + k_2 = \frac{2}{1+\beta+\beta^2}$
Simplifying further one can get $-\lambda\frac{\partial u}{\partial c_1^*} - \beta\lambda\frac{\partial u}{\partial c_2^*} + \frac{2\beta^2}{1+\beta+\beta^2}\lambda(1+2\frac{\partial f(a)}{\partial a}) = 0$. Ultimately, the expression (3) suppresses to

$$\frac{\partial \lambda}{\partial a}(k_1 + \beta k_2) = 0$$

Because $k_1 + \beta k_2 > 0$, $\frac{\partial \lambda}{\partial a}$ should be equal to 0, which means that $\frac{\partial f(a)}{\partial a} = \frac{1}{(2\beta^2+2\beta-1)}$ and the expression $\frac{1}{(2\beta^2+2\beta-1)} > 0$ for $\beta > \frac{\sqrt{3}-1}{2} \approx 0.37$

Appendix D

Let me consider a problem that is very similar to problems (1)-(2) with two differences. The first is that an agent has CARA utility function $u(c) = -exp(-x)$ and the second is that her assets are subject to an interest rate of R. I do these modifications for the following reasons. Under the quadratic utility function, the agent is ignorant about second-order uncertainty, so the sum of the consumer policy rule and saver policy rule coincides with a current income. On the other hand, the closed-form solution to the optimization problem with CARA utility exists only for income with a Gaussian distribution, which does not have the low bound $\underline{y}$. Thus, one cannot impose constraints on the problem (1)-(2) and no switching occurs. That is why I present two separate toy models for switching and the sun of consumption and savings policy rules. In the full infinite-horizon model with CRRA utility functions, which could be solved only numerically, both features will be present.

The second difference is the interest rate R that is imposed on the agent's assets, which she has access to in the last period only (due to infrequent assessment of assets). Thus, the interest rate is accumulated over the two periods. This fact drives the difference between

the policy rules of consumption targeters and savings targeters.

To remind you, an agent lives for 3 periods and creates a plan for all periods after observing income state μ. She does not update her plan as time goes on and does not violate it. The agent does not know her income y_2 and y_3 in the second and third periods. I define savings as $s \equiv y - c$.

The consumption plan is determined by the following optimization problem:

$$V^c = \max_{c_1,c_2} u(c_1) + u(c_2) + E_{y_2,y_3}[u(R^2[a + (\mu - c_1)] + R(y_2 - c_2) + y_3)] \tag{16}$$

The savings plan is determined by the following optimization problem:

$$V^s = \max_{s_1,s_2} u(\mu - s_1) + E_{y_2,y_3}[u(y_2 - s_2) + u(R^2[a + s_1] + Rs_2 + y_3)] \tag{17}$$

I have solved the model for the CARA utility case: $u(x) = -exp(-x)$ and under the assumption of Gaussian distributed income $y \sim N(\mu, \sigma^2)$. The unconstrained optimal solutions have the following form. Consumer optimally chooses:

$$c_1^* = \mu + \frac{R^2 a}{1 + R + R^2} - \frac{(1 + R^2)\sigma^2}{2(1 + R + R^2)}$$

Saver optimally chooses:

$$s_1^* = \frac{-R^2 a}{1 + R + R^2} + \frac{(1 + R\sigma)}{2(1 + R + R^2)}$$

Given these expressions, MPC out of assets for consumption targeters and savings targeters takes the form:

$$MPC^{consumer} = MPC^{saver} = \frac{R^2}{1 + R + R^2}$$

This example is essentially static, as the agent chooses her type once and for her whole life. In this example, I create switching decisions exogenously. The agent is a savings targeter and is endowed μ and a. She lives for three periods, dies and resurrects, receives Δ to her assets, and becomes a consumption targeter with μ and $a + \Delta$. For this switcher MPC takes the form:

$$MPC^{switcher} = \frac{R^2}{1 + R + R^2} + \frac{(R - R^2)\sigma^2}{2\Delta(1 + R + R^2)}$$

One could see that for the switcher the MPC does not coincide with the first derivative of the policy function and contains a bias. This happens because of non-linearity in future uncertainty. Because of this non-linearity consumption targeters and savings targeters treat

the uncertainty differently and $c + s \neq y$. Indeed,

$$c_1^* + s_1^* = \mu + \frac{(R - R^2)\sigma^2}{2(1 + R + R^2)}$$

where the last term occurs exactly because of the nonlinearity. More specifically, the distribution of the sum of incomes from the second and third periods weighted by interest rate R is $Ry_2 + y_3 \sim N([1 + R]\mu, [1 + R^2]\sigma^2)$. Thus, if one "bundles the uncertainty" (as a consumption targeter does), her optimal solution will contain σ^2 weighted by $[1 + R^2]$, while if she "splits the uncertainty" (as a savings targeter does), σ^2 in her optimal solution will be weighted by $[1 + R]$. This fact drives the gap between $c_1^* + s_1^*$ and μ and $c_1^* + s_1^* \neq \mu$ even though μ is known at the moment of choosing c_1^* and s_1^*. In other words, a consumption targeter sacrifices her consumption due to precautionary motives more than a savings targeter increases her savings because a consumption targeter mentally bundles uncertainty and perceives it to be higher than a savings targeter does.